

# Assignment 11: File system

**Due:** Friday, April 29, 10pm. Pre-submission due on **Tuesday, April 19, 10pm**. See [Deliverables](#).

**Starter code:** See [Assignment 11 on Canvas](#) for the Github Classroom link.

**Submission:** This is a **pair assignment**, but you can work alone, if you so choose.

Submit the contents of your repository via Gradescope. See [Deliverables](#) below for what to submit. If you are working with a partner, do not forget to include their name with the submission.

There will be no autograder for this assignment ahead of the deadline. Read the requirements and run tests locally.

Note: Clone and start studying the starter code as soon as possible. This assignment will likely require more programming effort than previous assignments.

## A File System

In this assignment you will build a [FUSE](#) filesystem driver that will let you mount a 1MB disk image (data file) as a filesystem.

### Step 1: Install FUSE

For this assignment you will need to use your local Ubuntu VM (or another local modern Linux). You'll need to install the following packages:

- `libfuse-dev`
- `libbsd-dev`
- `pkg-config`

Running

```
$ sudo apt-get install libfuse-dev libbsd-dev pkg-config
```

should do the trick.

Make sure your working directory is a proper Linux filesystem, not a remote-mounted Windows or Mac directory, i.e., do not use `/vagrant`.

After completing this step, go to [Pre-submission](#) and submit the repository as requested.

## Step 2: Implement a basic filesystem

You should extend the provided starter code so that it lets you do the following:

- Create files
- List the files in the filesystem root directory (where you mounted it)
- Write to small files (under 4k)
- Read from small files (under 4k)
- Rename files
- Delete files

You will need to extend the functionality in `nufs.c`, which only provides a simulated filesystem to begin with. This will require that you come up with a structure for how the file system will store data in its 1MB “disk”. See the [filesystem slides](#) and [OSTEP, Chapter 40](#) for inspiration.

We have provided some helper code in the `helpers/` directory. You can use it if you want, but you don’t have to. However, `blocks.{c,h}` and `bitmap.{c,h}` might save you some time as these implement block manipulation over a binary disk image. Feel free to extend the functionality if needed.

Some additional headers that might be useful are provided in the `hints` directory. These are just some data definitions and function prototypes to serve as an inspiration. Again, feel free to implement your own structs and abstractions.

## Step 3: Directories

In this step, implement support for arbitrarily nested directories. The filesystem should support the following operations on directories:

- Creation (`mkdir`)
- Renaming (`rename`)
- Listing the contents of directories (`readdir`)
- Deleting (`rmdir`)
- Creating files contained in directories, moving files between directories

## Step 4: “Arbitrarily” large files

Extend the filesystem to support files larger than 4K. The files must fit into the free blocks on disk. This must include proper allocation and deallocation as the file grows or shrinks. You should be able to support files with size of at least 40K.

## Deliverables

### Pre-submission

After doing [Step 1](#), that is, after cloning the repository and installing FUSE,

1. Execute the command

```
$ pkg-config --modversion fuse &> fuse_version
```
2. Commit the file `fuse_version` to your repo
3. Submit your repository to Gradescope under **Assignment 11: Pre-submission**

### Main submission

Modify the starter code to implement the requested functionality.

Commit the code to your repository. Do not include any executables, `.o` files, or other binary, temporary, or hidden files (unless they were part of the starter code). Do not include any disk images.

Once you are done, remember to submit your solution to Gradescope and do not forget to include your partner.

## Provided Makefile and Tests

The provided [Makefile](#) should simplify your development cycle. It provides the following targets:

- `make nufs` - compile the `nufs` binary. This binary can be run manually as follows:

```
$ ./nufs [FUSE_OPTIONS] mount_point disk_image
```
- `make mount` - mount a filesystem (using `data.nufs` as the image) under `mnt/` in the current directory
- `make unmount` - unmount the filesystem
- `make test` - run some tests on your implementation. This is a subset of tests we will run on your submission. It should give you an idea whether you are on the right path.
- `make gdb` - same as `make mount`, but run the filesystem in GDB for debugging
- `make clean` - remove executables and object files, as well as test logs and the `data.nufs`.

## Rubric

The grade is broken down into three categories:

- 5% Completing the **pre-submission** by Tuesday, April 19, 10pm
- 55% Basic functionality and FS design
  - Based on automatic and manual testing
  - Does the filesystem correctly and efficiently implement the requested functionality?
  - Do operations complete in a reasonable time? We put a 30s timeout on most test cases.
- 8% Directory functionality
- 7% Arbitrarily large files
- 25% Style
  - Via manual code review
  - Basics: meaningful purpose statements; explanation of arguments and return values
  - Explicitly stated assumptions
  - Correct use of types (e.g., not assigning -1 to an unsigned)
  - Short, understandable functions (generally, < 50 lines)
  - Consistent indentation and use of whitespace
  - Explanatory comments for complex blocks of code
  - No extra binaries (.o, executable files, etc.) or superfluous files committed to your repo

### Hints & Tips

- There are no man pages for FUSE. Instead, the documentation is in the header file: `/usr/include/fuse/fuse.h`
- The sources for `libfuse` contains a few further [examples](#). Start with `hello.c`.
- The basic development / testing strategy for this assignment is to run your program (e.g., using `make mount`) in one terminal window and try file system operations on the mounted filesystem in another separate terminal window.
- Read the manual pages for the system calls you're implementing.
- To return an error from a FUSE callback, you return it as a negative number (e.g. return `-ENOENT`). **Some things don't work if you don't return the right error codes.**
- Read and write, on success, return the number of bytes they actually read or wrote.
- You need to implement `getattr` early and make sure it's correct. Nothing works without `getattr`. The modes for the root directory and `hello.txt` in the starter code are good default values for directories and files respectively.
- The functions `dirname` and `basename` exist, but may mutate their argument.
- [https://www.cs.hmc.edu/~geoff/classes/hmc.cs135.201109/homework/fuse/fuse\\_doc.html](https://www.cs.hmc.edu/~geoff/classes/hmc.cs135.201109/homework/fuse/fuse_doc.html)

## Changelog

04/15/2022

- Remove a sentence that made it seem as if implementing deletion was optional.